

LoRa for the Internet of Things

Martin Bor
Department of Computing and
Communications
Lancaster University
m.bor@lancaster.ac.uk

John Vidler
Department of Computing and
Communications
Lancaster University
j.vidler@lancaster.ac.uk

Utz Roedig
Department of Computing and
Communications
Lancaster University
u.roedig@lancaster.ac.uk

Abstract

New transceiver technologies have emerged which enable power efficient communication over very long distances. Examples of such Low-Power Wide-Area Network (LPWAN) technologies are LoRa, Sigfox and Weightless. A typical application scenario for these technologies is city wide meter reading collection where devices send readings at very low frequency over a long distance to a data concentrator (one-hop networks). We argue that these transceivers are potentially very useful to construct more generic Internet of Things (IoT) networks incorporating multi-hop bi-directional communication enabling sensing and actuation. Furthermore, these transceivers have interesting features not available with more traditional transceivers used for IoT networks which enable construction of novel protocol elements.

In this paper we present a performance and capability analysis of a currently available LoRa transceiver. We describe its features and then demonstrate how such transceiver can be put to use efficiently in a wide-area application scenario. In particular we demonstrate how unique features such as concurrent non-destructive transmissions and carrier detection can be employed. Our deployment experiment demonstrates that 6 LoRa nodes can form a network covering 1.5 ha in a built up environment, achieving a potential lifetime of 2 year on 2 AA batteries and delivering data within 5 s and reliability of 80%.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

Keywords

LoRa, IoT, Medium Access Control (MAC)

1 Introduction

Recently new transceiver technologies have emerged which enable power efficient communication over very long

distances. Examples of such LPWAN technologies are LoRa [1], Sigfox [2] and Weightless [3]. These new transceiver types target applications where thousands of devices are used in a large geographic area to collect sensor readings. A typical application is the collection of meter readings in a city. These systems are used in a setup where simple devices send data in one hop to powerful receiver which then forwards data over a fixed wired infrastructure to a data collection point.

We argue that these transceivers are potentially very useful to construct more generic IoT networks incorporating multi-hop bi-directional communications enabling sensing and actuation. The transceivers have the ability to communicate over large distances on a small energy budget which would enable us to build more efficient IoT infrastructures than currently possible. For example, commonly used ZigBee transceivers such as the Chipcon CC2420 cover a communication range of 20 m using 84.5 μ J (40 byte message) in a built up environment. A LoRa Semtech SX1272 transceiver can cover a distance of 150 m using 86.5 mJ in the same environment.

Besides improved communication range, the transceivers have unique features stemming from the used modulation schemes. Thus, it is not efficient to simply use these transceivers with existing Medium Access Control (MAC) protocols and routing mechanisms that have emerged in the IoT domain. When constructing a network using these transceivers their specific capabilities should be taken into account to maximise performance in terms of communication and minimise energy consumption.

In this paper we investigate LoRa as technology for building generic IoT networks. We investigate the communication capability of the LoRa Semtech SX1272 transceiver and its energy consumption patterns. We analyse in detail unique communication features offered by the transceiver. A detailed analysis of these features is essential as they should be exploited when constructing communication protocols on top of this communication technology. Finally we construct an example communication protocol for the LoRa physical layer which enables wide-area multi-hop data collection and actuation without existing backbone infrastructure. Our deployment experiment demonstrates that 6 LoRa nodes can form a network covering 1.5 ha in a built up environment, achieving a potential lifetime of 2 years on 2 AA batteries and delivering data within 5 s and reliability of 80%.

The specific contributions of this paper are:

- *LoRa Feature Evaluation:* We describe and analyse LoRa specific physical layer capabilities. This includes channel separation using different Spreading Factor (SF), behaviour regarding non-destructive concurrent transmission; Clear Channel Assessment (CCA) in form of Carrier Activity Detection (CAD).
- *LoRaBlink:* We describe a protocol (LoRaBlink) on top of the LoRa's physical layer which exploits LoRa's unique features. The protocol enables energy efficient wide-area multi-hop data collection.
- *LoRaBlink Performance Evaluation:* We provide an evaluation of LoRaBlink in a small testbed using 6 nodes equipped with a Semtech SX1272 transceiver.

The next section describes our evaluation platform. Section 2 gives an overview of the LoRa Physical and MAC layer and Section 3 describes our LoRa evaluation platform. Section 4 gives an evaluation of LoRa specific features and Section 5 describes our proposed protocol LoRaBlink. Section 6 provides the evaluation of LoRaBlink and Section 7 concludes the paper. We do not provide a dedicated related work section as so far the research community has not investigated LoRa in depth and to the best of our knowledge no other research papers exist describing LoRa's capabilities and deployments.

2 LoRa

LoRa (**Long Range**) is a proprietary spread spectrum modulation technique by Semtech. It is a derivative of Chirp Spread Spectrum (CSS). The LoRa physical layer may be used with any MAC layer; however, LoRaWAN is the currently proposed MAC which operates a network in a simple star topology.

2.1 LoRaWAN

As LoRa is capable to transmit over very long distances it was decided that LoRaWAN only needs to support a star topology. Nodes transmit directly to a gateway which is powered and connected to a backbone infrastructure. Gateways are powerful devices with powerful radios capable to receive and decode multiple concurrent transmissions (up to 50). Three classes of node devices are defined: (1) Class A end-devices: The node transmits to the gateway when needed. After transmission the node opens a receive window to obtain queued messages from the gateway. (2) Class B end-devices with scheduled receive slots: The node behaves like a Class A node with additional receive windows at scheduled times. Gateway beacons are used for time synchronisation of end-devices. (3) Class C end-devices with maximal receive slots: these nodes are continuous listening which makes them unsuitable for battery powered operations.

In this paper we propose an alternative MAC for LoRa (LoRaBlink as described in Section 5) which enables multi-hop communication in a network of battery operated and duty-cycled devices. Although star networks with a powered and powerful gateway device are an option in some situations it does not cover all IoT application scenarios.

2.2 LoRa Physical Layer

LoRa is a physical layer specification based on CSS with integrated Forward Error Correction (FEC). Transmissions use a wide band to counter interference and to handle frequency offsets due to low cost crystals. A LoRa receiver can decode transmissions 19.5 dB below the noise floor. Thus, very long communication distances can be bridged. LoRa key properties are: long range, high robustness, multipath resistance, Doppler resistance, low power. LoRa operates in the lower ISM bands (EU: 868 MHz and 433 MHz, USA: 915 MHz and 433 MHz).

A LoRa radio has four configuration parameters: carrier frequency, spreading factor, bandwidth and coding rate. The selection of these parameters determines energy consumption, transmission range and resilience to noise. In the following sections we use the Semtech SX1272 transceiver as reference point.

Carrier Frequency

Carrier Frequency (CF) is the centre frequency used for the transmission band. For the SX1272 it is in the range of 860 MHz to 1020 MHz, programmable in steps of 61 Hz. The alternative radio chip Semtech SX1276 allows adjustment from 137 MHz to 1020 MHz.

Spreading Factor

SF is the ratio between the symbol rate and chip rate. A higher spreading factor increases the Signal to Noise Ratio (SNR), and thus sensitivity and range, but also increases the air time of the packet. The number of chips per symbol is calculated as 2^{sf} . For example, with an SF of 12 (SF12) 4096 chips/symbol are used. Each increase in SF halves the transmission rate and, hence, doubles transmission duration and ultimately energy consumption. Spreading factor can be selected from 6 to 12. SF6, with the highest rate transmission, is a special case and requires special operations. For example, implicit headers are required. Radio communications with different SF are orthogonal to each other and network separation using different SF is possible.

Bandwidth

Bandwidth (BW) is the range of frequencies in the transmission band. Higher BW gives a higher data rate (thus shorter time on air), but a lower sensitivity (due to integration of additional noise). A lower BW gives a higher sensitivity, but a lower data rate. Lower BW also requires more accurate crystals (less ppm). Data is sent out at a chip rate equal to the bandwidth. So, a bandwidth of 125 kHz corresponds to a chip rate of 125 kcps. The SX1272 has three programmable bandwidth settings: 500 kHz, 250 kHz and 125 kHz. The Semtech SX1272 can be programmed in the range of 7.8 kHz to 500 kHz, though bandwidths lower than 62.5 kHz requires a temperature compensated crystal oscillator (TCXO).

Coding Rate

Coding Rate (CR) is the FEC rate used by the LoRa modem and offers protection against bursts of interference. A higher CR offers more protection, but increases time on air. Radios with different CR (and same CF/SF/BW), can still communicate with each other. CR of the payload is stored in the header of the packet, which is always encoded at 4/8.

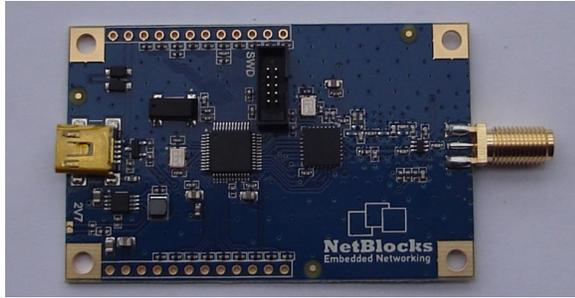


Figure 1. NetBlocks XRange SX1272 LoRa RF module.

2.3 LoRa Characteristics

Using a LoRa radio in a sensor network has some interesting aspects. First, since the range is relatively large (hundreds of meter indoors, kilometres outdoors), networks can span large areas without routing over many hops. In many cases one hop from every node to the sink is feasible. Secondly, transmission on the same carrier frequency, but with different spreading factor, are orthogonal. This creates the opportunity of dividing the channel in virtual subchannels. Thirdly, when transmissions occur at the same time with the same parameters, the strongest transmission will be received with high probability, i.e. concurrent transmissions are non-destructive even when their contents is different. This feature is exploited by LoRaWAN where all gateways broadcast beacons at the same time (tight clock synchronisation via GPS) and an end device is able to demodulate the strongest beacon.

3 LoRa Experimental Device

For our studies we use the XRange device from NetBlocks¹ as shown in Figure 1. The device comprises a SX1272 LoRa transceiver and a low-power STM32L151 ARM micro-controller. We use GCC ARM and the LoRa radio driver and runtime environment derived from the IBM LMIC (LoRaWAN in C)². The runtime environment, code used for our experiments and the communication protocol presented in Section 5 are available at <http://www.lancaster.ac.uk/scc/sites/lora/>.

The energy consumption of the system is for most application cases dominated by energy cost for communications. Energy consumption for transmission, reception, listening and CAD must be distinguished. The energy consumption in these states depends on selected SF and BW. Also, selected communication parameters will influence transmission times of packets and ultimately energy consumption.

To give an example we assume SF12, BW125, CR4/5, and TX power 17 dBm (An energy hungry setting allowing very long ranges which was used in our experimental evaluation discussed later). A transmission of a packet with 10 B payload and 12.25 symbols preamble has a transmission duration of 991.23 ms. Transmitting such message will cost 214 mJ. Reception of this message will cost 25.7 mJ and performing a CAD will cost 1.23 mJ.

¹<http://www.netblocks.eu/>

²<http://www.research.ibm.com/labs/zurich/ics/lrsc/lmic.html>

If we assume a system where the above message is transmitted every 15 min and we assume a battery capacity of 2 typical AA batteries of 5400 mAh the node will have a lifetime of 6.2 years.

If we assume a node only carries out a CAD every 5 s to check for an incoming message the node will have a lifetime of 6.0 years (assuming again 5400 mAh battery capacity).

4 LoRa Feature Evaluation

LoRa has interesting features aside the increased communication range that should be taken into account when constructing network protocols. For example, channel separation using different SF is possible, concurrent non-destructive transmissions are possible and carrier detection via CAD is provided. However, from available documentation the performance and ability of these features is not clear. Therefore we carry out a series of experiments to evaluate these provided features.

4.1 Spreading Factors

Different spreading factors are claimed to be orthogonal to each other. Thus, construction of virtual channels on the same carrier frequency is possible (Code Division Multiple Access (CDMA)).

We evaluate how well this separation works using a simple experiment setup. A transmitter is set to continuously send a 40 B packet with a fixed SF. A receiver set to the same SF is used to receive the transmissions. A second transmitter is used to transmit continuously and sequentially using all other SF to the same receiver.

Findings

All transmissions where sender and receiver use the same SF are received correctly. None of the transmissions emitted by the second node using a different SF are received. This result suggests that channel separation using SF works perfectly.

However, as we will show in Section 4.3 this is only partially true. When using CAD to detect an incoming transmission a signal using the wrong SF may be detected as valid transmission even though it cannot be decoded. This is important as the false detection rate has a negative impact on energy efficiency of a protocol.

4.2 Concurrent Transmissions

In LoRa concurrent transmissions are claimed to be non-destructive and such feature is very valuable for protocol design. Well-timed cooperative transmissions have been used in Glossy [6]. In Glossy the same message is transmitted accurately timed by multiple nodes allowing correct reception. A-MAC [5] and also Whitehouse et al. [7] make use of the capture effect. Here multiple different messages are transmitted concurrently and depending on power levels and timing one of the concurrently transmitted messages can be received.

We set up an experiment to understand the exact conditions in which this effect is present in LoRa. We use a receiver, one 'weak' transmitter and one 'strong' transmitter (1 dBm difference). Both transmitters send a packet with explicit header and CRC. The strong transmitter varied the transmission time offset relative to the weak transmitter. From being one packet (airtime) early to being one packet

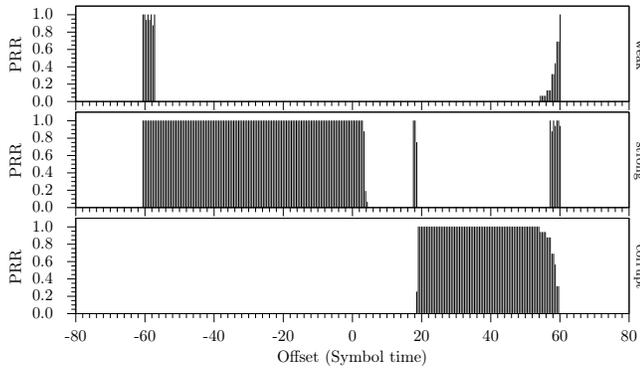


Figure 2. Example collision result. Spreading factor 11, bandwidth 125 kHz. X-axis shows the transmission offset relative to the weak node in symbol time, Y-axis shows the packet reception rate.

(airtime) late. For each offset, 16 32-byte packets were transmitted using first identical packet payloads and subsequently different payloads. We also run the experiment with all combinations of SF and BW.

The experiment results are shown in Figure 2 for SF12 and BW 125 kHz. The Y-axis represents the packet reception rate. The X-axis represents the transmission offset relative to the weak node in symbol time. The top bar shows packets received from the weak transmitter at the receiver; the middle bar shows packets received from the strong transmitter at the receiver. The bottom bar shows when packets were received from either transmitter but deemed corrupt (Cyclic Redundancy Check (CRC) failure). We did notice that about 1 in 6000 packets was corrupted, but did not fail the CRC. Often these packets had 1 bit corrupted.

Results for other SF and BW combinations are very similar. Also, transmitting the same packet payload or a different payload does not change the obtained results significantly.

As can be seen, the strong transmitter is successfully decoded if it transmits not later than 3 symbol periods after the weak transmitter started (symbol period is $T_{sym} = 2^{SF}/BW$). If the weak transmitter starts later than 3 symbol periods no transmission is received (or corrupted data is received).

Although the packet takes 60.25 symbol periods to transmit, both nodes can be received at an offset of -57 symbol periods or more. The tail of the strong node does destroy the initial preamble of the weak node, but as long as at most 3 symbols are destroyed, the weak packet can also be successfully received. This relationship is not symmetrical, as at an offset of $+57$ symbol periods, the weak node's tail (CRC) gets destroyed, invalidating a packet which may have been correctly received. It is only that at an offset of $+60$ symbol periods or more, both packets gets received perfectly.

We also experimented with two transmitters set to the same transmit power. In this case either of the two is perceived as stronger and the above described behaviour applies (although the role of stronger/weaker transmitter may alternate with each transmission making it difficult to conduct experiments and describe results).

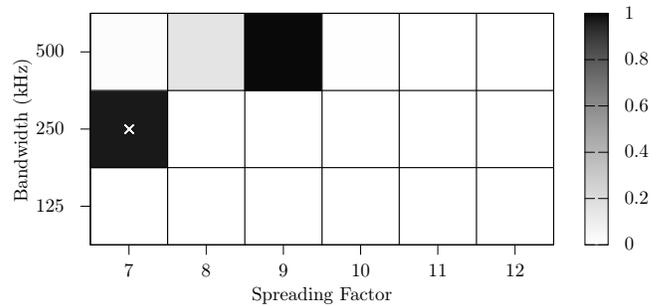


Figure 3. Carrier detection ratios for a transmitter sending at spreading factor 7 and bandwidth 250 kHz, indicated by the white cross. Carriers were also detected by adjacent data rates.

Findings

One of two concurrent transmission can be received with very high probability if both transmissions do not have an offset of more than 3 symbol periods. This translates to a duration between $768 \mu s$ and $98.3 ms$, depending on the SF and BW. Synchronisation of nodes within these bounds is relatively easy to achieve and therefore protocols making use of this feature can easily be implemented with LoRa.

4.3 Carrier Activity Detection

Transceivers normally provide a CCA interface to detect an occupied channel. The CCA is used in communication protocols to decide if packets can be transmitted and to decide if the radio must be kept active to receive a message. In particular for power constrained nodes it is important to have an accurate and fast CCA mechanism as this enables implementation of power-efficient duty cycling. Nodes perform periodic short CCA checks and only power the receiver for longer if a transmission is detected.

LoRa transceivers do not provide a classical CCA interface based on an Received Signal Strength Indicator (RSSI) threshold to detect an occupied channel. LoRa can receive transmissions with a signal strength that is below the noise floor and, consequently, an RSSI threshold check will not reveal an occupied channel. LoRa radios provide therefore a CAD mode to detect a present preamble.

The CAD process takes approximately 2 symbol periods, and only requires the radio on for about 1 symbol period (The exact CAD duration is calculated as sum of $(32/BW + 2^{SF}/BW)$ seconds in RX mode and $(SF \cdot 2^{SF})/(1750 \times 10^3)$ seconds of processing). The processing phase requires about half of the energy required in receive mode (around 6 mA depending on the SF/BW). After receiving a 'CAD detected', a transceiver can be switched in RX mode to receive the ongoing transmission if required.

We set up an experiment to test the reliability of CAD. A detector node starts the CAD process on a regular interval (every 100 ms) and records whether it has detected a carrier or not. After 300 samples, it switches the SF/BW combination and repeats the process. A transmitter node is programmed to continuously send out preambles at a fixed SF/BW combination. The experiment is repeated with different transmitter SF/BW combinations.

The results for a transmitter using SF7 and a BW of 250 kHz are shown in Figure 3. Results for a transmitter using different SF/BW combinations are similar. When transmitter and receiver use the same SF/BW combination the worst detection rate was measured at 97% (for SF7 and a BW of 250 kHz). However, the CAD process also detects carriers in SF/BW combinations different from the combination the transmitter is using (up to 99% detections for SF9 and a BW of 500 kHz). This happens for data rates that are adjacent to the current data rate. When no transmitter is active, the detector had a false positive rate of 0.092%.

Findings

CAD can only detect channel occupancy while a preamble is transmitted. The detection probability is high (above 97%) and false positives are low (0.092%). However, if multiple LoRa networks are active on different SF/BW combinations false positives can be very high (depending on SF/BW ratios). When using multiple SF/BW combinations in the same network (or when constructing multiple networks separated by SF/BW) the choice of combinations is important when using CAD.

5 LoRaBlink

In this section we describe LoRaBlink, an IoT protocol for LoRa transceivers. The protocol is designed to support reliable and energy efficient multi-hop communication. It is also designed to support low-latency bidirectional communication. Thus, the described protocol is different to defined protocols for LoRa as given in the LoRaWAN 1.0 specification. The protocol relies on features and building blocks as described in the previous section.

5.1 Protocol Aims

LoRaBlink aims to address a number of aspects necessary for deployment of IoT applications and which are not covered by currently defined LoRa protocols. These are:

- **Multi-Hop:** The protocol should support communication over multiple hops.
- **Low-Energy:** Nodes should be able to duty-cycle to conserve energy and enable battery powered operations over long time spans.
- **Resilience:** The protocol should be resilient and enable high message delivery probability.
- **Low-Latency:** The protocol should enable low-latency communication.

Further to these requirements we also make the assumption that the network has a low density, low traffic volume and contains a limited number of nodes. We also assume that a single sink is used for communication and that communication is between the sink and the nodes.

A vast number of protocols exist to implement these requirements [4]. However, none of the available options is particularly designed to make use of LoRa specific features such as the ability to receive one message out of a pool of concurrent transmissions.

5.2 Protocol Operations

The protocol integrates MAC and routing in a single simple protocol. Time synchronisation among nodes is used to define slotted channel access. Nodes transmit concurrently

within slots and properties of the LoRa physical layer ensure that one of the concurrent transmissions is received. Messages are distributed from the sink to nodes using flooding. Messages from nodes to the sink use a directed flooding approach. The result is a very simple but robust protocol which covers the set requirements.

Figure 4 shows an operation example of LoRaBlink in a network containing 3 nodes and a sink. Node 1 and 2 are in communication range of the sink. Node 3 cannot be directly reached by the sink but is in range of node 1 and 2.

Each node powering up will remain in listen mode until a beacon is received. Beacons are used for time synchronisation and mark the start of an epoch. Each epoch contains N slots. The first N_B slots of an epoch are used for beacon transmissions. A beacon message contains the hop distance to the sink and upon receiving a beacon a node will transmit its own beacon according to its distance to the sink. A node will aim to select its position based on minimal distance to the sink. In the example in Figure 4 the sink transmits a beacon received by node 1 and node 2. Both nodes use the beacon to determine epoch start and their distance to the sink (1 hop). In the next beacon slot node 1 and 2 transmit their beacon concurrently. Due to properties of the LoRa physical layer either one of these (depending on transmission time difference and perceived signal strength at node 3) is received at node 3. Node 3 updates its hop distance to the sink as 2 and transmits its own beacon in the next beacon slot. This beacon is received by node 2 (we assume node 1 would not receive) which discards the message as its hop count is less than 2. The number of beacon slots N_B determines the maximum depth the network can have.

Following the beacon slots are N_D data slots. A node that has data to transmit selects the next available data slot and transmits. After transmission a node listens for an Acknowledgement (ACK) (an optional protocol feature; the ACK is not shown in Figure 4). Two nodes may transmit in the same slot with the result that at least one message is decoded by one receiver, with a chance that two different nodes in the network decode one of each transmission. If a node has a lower hop count to the sink than the source node it will relay the message in the next slot. Multiple nodes may forward which introduces redundancy. ACK messages may also collide but a receiver will always be able to decode one of multiple ACK correctly. In Figure 4 node 3 generates a data message. The message is received by node 2 and 1 which then forward the message simultaneously in the next slot. The sink will be able to decode one of the two transmissions. Data travelling from the sink to a node will use the same mechanism as used for beacon distribution. If the sink has to send non-delay sensitive information to nodes in the network it can be delayed and included in beacon messages for distribution.

5.3 Node Lifetime

To improve energy consumption of the system beacon messages are sent infrequent (a long epoch is used) and the CAD is used within slots to detect incoming transmissions. Infrequent beacon transmission is possible as tight time synchronisation in the network is not necessary.

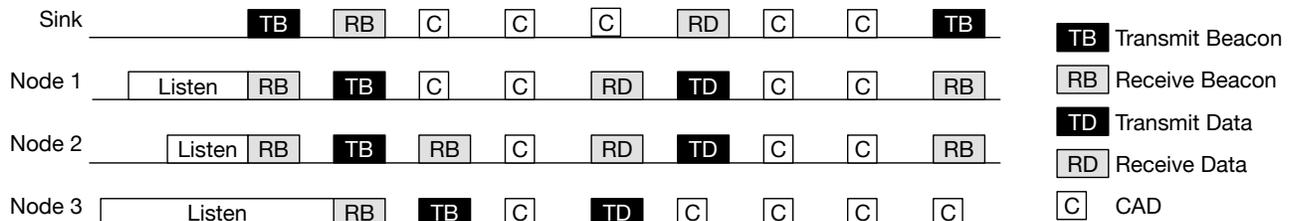


Figure 4. LoRaBlink: Protocol example using a 4 node network.

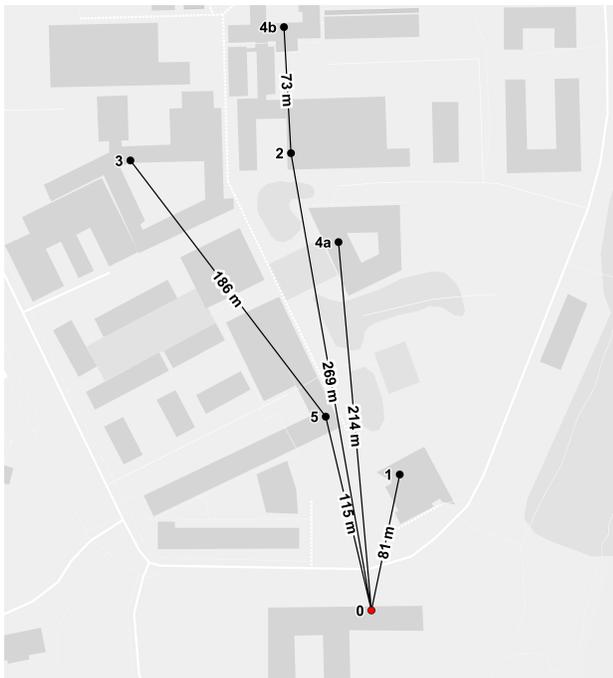


Figure 5. LoRaBlink: Map of a small scale deployment. Lines are routes between nodes, with distance in meters.

Epoch length and N_B and N_D determine energy consumption and data transport delay in the network.

While not the most energy efficient protocol, the additional range may be a benefit to low-node-density deployments, requiring far fewer forwarding nodes to cover the same area. The transceiver settings described in Section 3 and an epoch length of 15 min with $N_B = 3$ and $N_D = 177$ (5 second slot) we obtain a maximum node lifetime of 2 years with two AA batteries (5400 mAh) – assuming that one beacon is transmitted and two are received and all other slots in the epoch contain one CAD.

6 LoRaBlink Evaluation

We deployed a 6 node network on the university campus as shown in Figure 5. Nodes are deployed in buildings across campus on the ground floor within buildings approximately 1.5 m above the floor. The sink node is located in the third floor on a windowsill. Node 4 was first deployed at position 4a and was moved to position 4b to create a larger network. Nodes have to communicate through several buildings and structures.

We use in the experiment SF12 and BW125 and a TX

power of 17 dBm. The epoch length was set to 5 min with $N_B = 5$ and $N_D = 55$ (slots every 5 seconds). Nodes are set to transmit a data packet (10 byte) randomly within one slot of each epoch. In this experiment we did not use CAD and instead implemented a listen period of 50 symbols in each slot. This was done to avoid packet losses due to CAD and to evaluate data delivery of LoRaBlink on its own.

In our evaluation packets from all nodes were delivered with a reliability of 80% over a duration of 2.3 h. Node 4 delivered messages for the first half of the experiment from the position marked 4a and later from position 4b. Node 3 and 4b delivered messages via one hop while all other nodes were able to directly communicate with the sink node. Messages transmitted by node 3 and 4b are relayed by multiple nodes (node 5, node 2 and node 4a for node 3) in the same slot.

The experiment shows that LoRaBlink can deliver messages reliably over large distance in a challenging multi-hop environment (structures in the communication path). The experiment also shows that using concurrent transmissions is feasible.

7 Conclusion

LoRa radios are capable of longer communication ranges than commonly used IoT radios while still being energy efficient. In addition, these radios provide interesting features such as non-destructive concurrent transmissions. As we have shown, LoRa radios can be used in more general network layouts than the one used by LoRaWAN. Thus, we believe that LoRa transceivers provide an interesting option for building general IoT applications.

8 References

- [1] LoRa. <https://www.lora-alliance.org>. Accessed: 2015-11-07.
- [2] Sigfox. <http://www.sigfox.com>. Accessed: 2015-11-07.
- [3] Weightless open standard. <http://www.weightless.org>. Accessed: 2015-11-07.
- [4] A. Bachir, M. Dohler, T. Watteyne, and K. K. Leung. Mac essentials for wireless sensor networks. *Communications Surveys & Tutorials, IEEE*, 12(2):222–248, 2010.
- [5] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis. Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, SenSys '10*, pages 1–14, New York, NY, USA, 2010. ACM.
- [6] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time synchronization with glossy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 73–84, April 2011.
- [7] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Embedded Networked Sensors, 2005. EmNetS-II. The Second IEEE Workshop on*, pages 45–52, May 2005.